Design of Access Control Methods for Protecting the Confidentiality of Patient Information in Networked Systems

Jesse W. Bowen, Ph.D., M.S.E.E.[1], J. Craig Klimczak, D.V.M., M.S.[1,2],
Michael Ruiz, R.N., M.B.A.[1], and Mike Barnes, M.D.[1]
[1]Integrated Technology Services and [2]Health Services Management,
University of Missouri-Columbia School of Medicine
Columbia, Missouri

ABSTRACT

Public awareness of the potential for violation of personal privacy in clinical information systems is increasing. Much of this increase can be attributed to the popularity and publicity of the World Wide Web. Nightly news reports of intruder break-ins and flaws in Internet software security have stimulated public interest in the security of clinical information systems available over the web. As part of the development of systems designed to provide clinical narratives to physicians over the Internet, we are exploring designs that provide additional protection and security to these systems. Specifically, we are developing and testing automated access control measures based on provider-patient relationships for controlling access to personally identifiable patient information.

INTRODUCTION

Concerns about protecting the confidentiality of patient and medical data has recently emerged as a major challenge facing designers of medical information systems (MIS). Growing public awareness of how electronic information is collected and stored, coupled with increasing sensitivity to the potential damage that can result from public disclosure of some medical information, has prompted introduction of federal legislation [1], public hearings [2], and a National Research Council report [3] on the topic. At present, no compelling consensus has formed to define the rights of patients with respect to their medical data stored electronically, complicating the design task for MIS security mechanisms. Yet it is probable that in the near future legislative or regulatory action will assert the rights of individuals to the protection and control of medical information in a manner similar to principles embodied in the Fair Information Practices [4] and recommendations of the Computer-based Patient Record Institute (CPRI) [5]. It is apparent from reports in the informatics literature that there are no methods that address such security concerns are clearly superior to others; security systems depend to a great extent on the goals of the system designers and the institutional policies they seek to implement.

Our institution has developed a distributed client/server system that provides access to transcribed clinical narratives, using a World Wide Web browser as client [6]. As part of a National Information Infrastructure contract [7], we are planning to provide access to referring physicians in the surrounding community, some of whom will not be employees of this institution. In an effort to address any special security problems that result from this service, and to prepare for the possibility of more stringent security requirements as policy initiatives develop, we have analyzed a variety of approaches to protecting patient confidentiality. In this report we will describe some of the MIS security features used by other institutions as reported in the literature and detail their advantages and shortcomings for protecting patient confidentiality in our own institution. We will then describe the related security measures we believe hold promise for effectively protecting medical information in our environment.

CURRENT APPROACHES

Security mechanisms designed to protect patient confidentiality generally rely on some combination of authentication, authorization, and auditing. *Authentication* refers to the means by which an information verifies the identity of a user, usually based on passwords or physical tokens. *Authorization* will denote access controls or other means used to provide specific information resources to a given user. *Auditing* will be used to describe processes for recording and reviewing a user's interaction with the system.

While few systems rely solely on only one of these three components, some security designs place

1091—8280/97/$5.00

primary emphasis on user authentication [8]. After authentication, users have access to all patient data in the system. In the cited report, any loss of patient privacy was deemed justified, on balance, to achieve the goal of maximal data sharing. More commonly, systems combine user authentication with some form of audit trail [9,10,11]. The retrospective information captured by recording user transactions is designed to determine responsibility in the event of confidentiality breaches, or to serve as a deterrent to inappropriate use. For the latter purpose, confidentiality agreements with users are usually required. At least one system [9] also has periodic reminder screens to reinforce user appreciation of the auditing process. While all installations implement some form of authentication, relatively few institutions [12,13] use access controls. In the recent NRC report [3], only one of the six sites investigated had access controls. The NRC report endorsed robust authentication and auditing, but it also recommended the further development and wider use of authorization procedures.

The great variety in approaches to protecting patient confidentiality derives in part from differences in the institutional policies supported by the security systems. Systems designed to promote the greatest possible access to highly trusted and responsible users tend to rely more heavily on user authentication. Institutions that desire to deliver information to a wider variety of users often implement some form of access control. There is general recognition of the principle that access to information and its security are inversely proportional. As a result, it is necessary to strike a balance that supplies useful information to health care providers while minimizing risks to confidentiality. The criterion cited most often to make such a decision is whether a given provider has a "need-to-know" the information in question. The crucial aspect in protecting confidentiality may be viewed as the mechanism by which a security system makes the need-to-know judgement for a given unit of patient information. The following section examines how various security arrangements make this decision.

ANALYSIS OF SECURITY DESIGNS

Authentication

Protections based only on authentication steps have the advantage of providing wide access to data at relatively low implementation costs and demands on system performance. Such security designs make the assumption that any valid user of the system has the right to decide for themselves whether they have the need-to-know for available information. This approach usually involves users who are all physicians or other trusted persons deemed to possess a high degree of responsibility. Reasons given to support this assumption include:

- physicians have free access to paper-based medical records
- physicians are ultimately the responsible parties for health care decisions
- restrictions on access to records may compromise the quality of care
- physicians are already bound by patient confidentiality agreements

It may be useful to re-examine some of these assertions in light of differences between electronic medical records and their paper-based counterparts. For example, in most institutions it is not difficult for a physician to request access to any medical record. But there are practical difficulties associated with actually obtaining the physical chart. Thus "free access" to existing paper records is primarily theoretical and does not correspond exactly to the concept of free access to electronic versions of patient data. It may be possible to implement substantial controls on access to electronic patient records and still have a system that provides greater access to patient information than a paper-based system. The concept that quality of care may be compromised by restrictions on access does not account for the fact that restrictions need not be absolute. In fact, many systems supplement such restrictions with methods for overriding them in time-critical or emergency situations.

The idea that free access to information is justified because of the high degree of trust placed in physician users bears examination on two fronts. First, it can be anticipated that there will be finite rates of abusive access to confidential information by even the most trusted groups of users, although they may be very low. It can also be expected that abuses will increase as the number of system users increase and as the quantity of information provided by the system increases. A fundamental design decision must therefore be made to determine whether the frequency and severity of such abuse is acceptable to the institution in terms of possible legal or financial consequences. Second, most systems already allow some limitations even for the most trusted users. This may be in the form of special handling procedures for records of employees or others, such as celebrities, considered to have VIP status. At our institution, for example, it is possible for patients to request special handling of medical records to limit or prevent their access by

specific individuals, including physicians. Electronic medical record systems that depend only on user authentication will lack these forms of access control currently implemented by paper-based systems.

The most serious limitation of authentication-based systems is that they don't closer control for users who are not physicians or other highly trusted individuals. As the quantity and quality of electronic clinical information increases, it is to be expected that a widening population of health care workers will require access to specific types of data. Even though the risk of potential abuse may be small, the overall number of violations will increase with both the number of users and the amount of information available. As a result, many institutions have used auditing systems along with authentication to create indirect means for reducing abuses of confidentiality.

Auditing and Authentication

To support authentication and address shortcomings of authentication-only systems, many institutions record user transactions to create an audit trail. Provided that sufficient detail is captured, user education and the threat of sanctions can be used as a deterrent to abuse of information. Depending on the processing capabilities of the information system, auditing may appear transparent to the user. Audit trails then have the potential to decrease abuse without erecting functional barriers to access. Audit records are also useful for a variety of other purposes, such as system monitoring or to obtain data about patterns of access.

Auditing systems have some intrinsic disadvantages. They require additional system resources for monitoring, storing, and managing transaction records. The additional effort required to record the transaction information may compromise overall system response, and the decline increases as the recorded detail increases. Perhaps the most serious demand made by audit trails is the need to provide policies, personnel, and resources to review the resulting information. At one institution with a large number of users [12], about 100 MB of data was generated per month, roughly equivalent to 30,000 printed pages. Some of the work associated with review of this data may be alleviated by developing systems that identify and flag possible security violations as they occur.

The primary limitation of security systems based on authentication and auditing alone is that they are primarily retrospective. Such a system depends on users being aware of and sensitive to the consequences of abuse at some later time. Sanctions related to employment or revoking system privileges will only be effective for employees of the institution who wish to retain access to the system. Thus auditing and authentication alone will not prevent abuse by a valid system user who for some reason chooses to ignore such consequences.

Access Control

Systems that implement authorization procedures generally attempt to determine whether a given user has need-to-know for the requested information. There is great variety in the types of information used to formulate mechanisms for making this decision [13,14,15]. Some of the database dimensions that have been proposed for this use include attributes related to segments of records themselves, time periods for access, types of users, types of database interactions, the purpose for which the information will be used, and the need for information across patient populations as opposed to individual patient data [14,15]. A common paradigm is to develop a rights matrix with two or more dimensions, and assign an appropriate level of access for each matrix element. For example, an access matrix could be based on user roles and types of patient information. In this case, each role would have a defined level of access for each information type. The advantage of this approach is that access control may be specified along multiple dimensions with as much detail as desired. The concomitant disadvantage is that increasing complexity enlarges the effort required to establish and manage the rights matrix.

Despite the potential cost in resources, it is probable that authorization procedures to implement access control will be required as mandates for protection of patient confidentiality develop over the next several years. There are also legal precedents finding that merely instructing employees not to violate confidentiality is insufficient without adequate institutional policies and security structures to ensure compliance [16].

METHODS OF ACCESS CONTROL

Our institution has implemented the System for Text Archive and Retrieval (STAR) to provide free-text clinical narratives via an Internet browser client [6]. The initial security provisions for this system have been limited to password-based authentication and a simple access log. There are currently about 50 users with access to the system, consisting of physicians and system developers. We have been investigating

the implications for access control systems of providing STAR capabilities to referring physicians from the surrounding area who are not employees of this institution. Our desire is to develop methods for access control that can be implemented incrementally as system capabilities improve and institutional policies are formulated, yet would provide robust protections for patient confidentiality. The goal was to design an access control system that would be adaptive, in that it would be able to use system information derived from registration systems to make access control decisions. Such a system would initially combine a limited number of user roles with access permissions granted by checking for established relationships between providers and patients. Some of the components considered for inclusion in such an integrated authorization system are discussed below.

Provider-Patient Relationships

One way of ensuring that system users have access only to information for which they have a valid need-to-know is to require that there be a pre-existing and explicit relationship between a health care provider and the patient for whom information is sought. Such a provider-patient relationship (PPR) may be established in a variety of ways. Manually constructing a database of such relationships would be tedious for a large number of users. Another possibility would be to allow users to establish relationships themselves. This method would not prohibit inappropriate accesses by itself, but may be useful as a way of both warning system users of what is considered valid use of the system and to provide security alerts to aid interpretation and examination of audit logs. An efficient technique that retains strong access control would be to use patient registration information to generate PPR data.

We have examined the existing structure of STAR to determine whether currently stored information could be used to establish a valid PPR. STAR contains relational database tables that store information on patients, health care providers, and text documents. In addition, admission-discharge-transfer (ADT) data from the hospital registration system is captured through an HL-7 interface and used to populate a case cross-reference table with identifiers for patients and physicians associated with a case. During a recent five month testing period during which STAR was made available to 28 University Hospital physicians (none directly associated with the development team), there were 1018 requests for patient documents. Of these, 40 were from physicians who requested access to patients with whom they were associated in the case cross-reference table. This implies that, as currently configured, only about 4% of access requests would have been identified as possessing a valid PPR. This rate obviously must be much higher if ADT data is to be used effectively for automated generation of PPR data. There are several reasons for the lack of sufficient data linking providers and patients. As currently designed, STAR receives ADT information only from the main hospital registration system and not from outpatient clinics and an ancillary cancer treatment facility, both of which use registration systems from different vendors. Incomplete and inaccurate data, manually entered at the time of admission, has also been identified. Document profile data and the archived documents themselves are both valuable sources of data linking providers to patients, but neither is currently entered into the case cross-reference table.

Until information from registration systems can be obtained more systematically, manual entry of PPR data will be used. One of the first applications will be to provide access control for non-employee users. There will be two user roles, one for employee physicians and one for non-employee physicians, and a valid PPR will be required for users in the latter role to gain access to a patient's documents. It may also be desirable to use the PPR mechanism for employee users as a means to establish "negative" relationships, to prevent access to specific patient records. Concerns about the performance cost of querying and maintaining a large database table to store PPR data could be addressed by developing a cache system. A PPR cache could retain only a subset of most-requested or most-useful information that would include, for example, only in-patients, patients scheduled for imminent visits, and recently discharged patients. Such a cache would have the added benefit of allowing efficient identification of patients to present in a patient list tailored for specific users.

Flexible Barriers and Auditing Policies

As an adjunct to user roles and PPR data, auditing policies and procedures can provide an indirect means of access control. Despite the limitations noted above, for most users the knowledge that transactions will be recorded and reviewed should provide a powerful motivation to avoid abuse. We have considered audit policies that would supplement this goal. An obvious and relatiely easy task is to design an ongoing educational program that clarifies the institution's definition of appropriate use, describes the system's auditing capabilities, and specifies sanctions for

confidentiality violations. An additional technique would be to use "flexible" barriers as a way of reinforcing awareness of the auditing process and flagging user actions that require security review. Barriers could range from special screens that demand user interaction such as clicking a button before proceeding, up to a requirement that a password be entered to gain access to especially sensitive data. Audits of barrier events would provide one way to identify transactions with exceptional security interest.

Another technique that holds promise for decreasing the burden of reviewing audit data is to allow self-audits by patients. Either routinely or on request, lists of users who have requested access to records of a given patient could be generated and supplied to the patient for examination. It is possible that patients could identify inappropriate access in cases where institutional procedures would fail because of insufficient information about a patient's friends or acquaintances. This would also serve as an additional deterrent to casual access if this capability were advertised to users.

CONCLUSIONS

There are several reasons why proactive access controls are desirable additions to the essentially retroactive protections afforded by systems that rely primarily on authentication and auditing. As the capabilities of electronic clinical information systems improve, the challenge will be to provide effective protections for patient confidentiality that have a minimum impact on the quality of health care. To this end, it will be useful to search for accurate sources of information and combinations of security mechanisms that allow appropriate balance between the information needs of system users and the protection of patient confidentiality.

Acknowledgments

This work was supported in part by NLM Grant LM07089 and Contract LM63538. Dr. Bowen is supported by NLM Training Grant LM05415.

References

1. S. 1360, Medical Records Confidentiality Act of 1995; H.R. 435, Fair Health Information Practices Act of 1995; H.R. 3482, Medical Privacy in the Age of New Technologies Act of 1996.
2. Department of Health and Human Services National Committee on Vital and Health Statistics: Subcommittee on Health Data Needs, Standards, and Security Meetings. Federal Register, 1997; 62(6):1336-7.
3. National Research Council. For the record: protecting electronic health information. National Academy Press, 1997; pre-publication copy
4. Kluge, E-HW. Health information, the fair information principles and ethics. Methods Inform Med, 1994; 33:336-45.
5. Computer-based Patient Record Institute. Position paper: access to patient data. http://www.cpri.org/docs/access.html, 1996; Dec. 3
6. Klimczak CJ, Bopp K. Reengineering medical records with a text archive and retrieval system. HIMSS Proc, 1996; 3:63-76
7. Mitchell, JA. Expand the benfits of rural telemedicine services by linking health professionals in three small Missouri communities. NLM Contract LM63538, 1996.
8. Gardner, RM. Integrated computerized records provide improved quality of care with little loss of privacy. JAMIA 1994; 1(4):320-22.
9. Safran C, Rind D, Citroen M, Bakker AR, Slack WV, Bleich HL. Protection of confidentiality in the computer-based patient record. Clin Comp 1995; 12(3):187-92.
10. Wear PK, Skeens MA, Thorne C. Building security models for patient identifiable health information. HIMSS Proc, 1996; 3:237-53
11. Chute CG, Crowson DL, Buntrock JD. Medical information retrieval and WWW browsers at Mayo. JAMIA, 1995; suppl:903-7.
12. Barrows, RC Jr., Clayton PD. Privacy, confidentiality, and electronic medical records. JAMIA, 1996; 3(2);139-48.
13. Dargahi R, Classen SW, Bobroff RB, *et al.* The development of a data security model for the collaborative social and medical services system. JAMIA, 1994; suppl:349-53.
14. Brannigan VM. A framework for "need to know" authorizations in medical computer systems: responding to the constitutional requirements. JAMIA, 1994; suppl:392-96.
15. Henkind SJ, Orlowski JM, Skarulis PC. Application of a multilevel access model in the development of a security infrastructure for a clinical information system. 17th SCAMC Proc, 1993; 64-8.
16. Brannigan VM. Protection of patient data in multi-institutional medical computer networks: regulatory effectiveness analysis. 17th SCAMC, 1993; 266-70.